

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY) 01/20/2012		2. REPORT TYPE Final		3. DATES COVERED (From - To) 5/1/2010 - 4/30/2011		
4. TITLE AND SUBTITLE MURI: CARA: Cognitive Architecture for Reasoning about Adversaries				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER FA9550-06-1-0405		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Subrahmanian, V. S. Nau, Dana				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland College Park, MD 20742				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) USAF, AFRL 875 N. Randolph Street 4036 Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR/RSL		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-OSR-VA-TR-2012-0160		
12. DISTRIBUTION/AVAILABILITY STATEMENT A						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT We developed the CARA (Cultural Adversarial Reasoning Architecture) that provides a software platform using which analysts can reason about cultural groups, tribes, and terror groups. We developed the T-REX system to automatically extract information about violent events from news sources worldwide, as well as Stochastic Opponent Modeling Agents that allow us to reason about adversarial behavior. We developed two algorithms to make forecasts about adversarial behavior. We developed game-theoretical approaches to reason about group behavior. We tested out our approaches with real world data related to Hamas and Hezbollah and made forecasts that were true even though the terror group in question knew about the forecasts.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON V.S. Subrahmanian	
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code) 301-405-6724	

I. Accomplishments

The goal of the “Cognitive Architecture for Reasoning about Adversaries” project (or CARA for short) was to develop the mathematical foundations, data structures, algorithms, prototype implementations, and validation methods to enable researchers, defense analysts, and policy makers to learn models of the behaviors of groups (irrespective of whether they are terrorist groups or ethnic groups or tribes or political parties or investor groups) and to effectively use those models to advance their mission.

CARA made significant advances – not just in scientific accomplishments that significantly extend the state of art, but also in significant outreach efforts to build a broad scientific community focused on human socio-cultural behavioral modeling.

I.A Scientific Accomplishments

In a nutshell, CARA made the following major scientific accomplishments:

- Proposed the first *advanced software architecture* that would identify the necessary components and communication methods needed to build an end-to-end cultural reasoning system;
- *Advanced the state of the art in information retrieval and natural language processing* so as to automatically collect information about events related to the construction of such models;
- *Developed the first syntax within which behavioral models of groups could be expressed;*
- *Developed the necessary mathematics and algorithms to use these models to automatically make forecasts about group behavior* together with methods to quantify the uncertainty inherent in such forecasts;
- *Developed algorithms to make forecasts of group behavior based on past behaviors;*
- *Successfully made forecasts about several real-world terror groups such as Hezbollah and Lashkar-e-Taiba on the basis of these technologies;*
- *Understand how human group decision making operates in reality* as opposed to theoretical models based on expectation maximization;

- *Developed evolutionary game-theoretic methods to understand the evolution of group cooperative behavior and its relationship to the behavior of individual agents;*
- *Developed methods to automatically distinguish between intentional actions of an agent and noise that tends to abound in real-world systems;*
- *Develop game-theoretic methods to learn the payoffs and outcomes of behaviors of other entities (and co-opt them into one's own behaviors);*

I.B Outreach Efforts

Computer technology is leading to sweeping changes in how we can reason about groups in diverse cultures. Examples include computer systems to aid researchers in gathering data about different cultural groups, learning the intensity of opinions that those groups have on various topics, building/extracting models of behavior of those groups, and continuously refining those behaviors through shared, multi-person, learning experiences. These developments are inherently cross-disciplinary, blending the behavioral and social sciences—fields such as political science, psychology, journalism, anthropology, and sociology—with technological fields such as computer science, computational linguistics, game theory, and operations research. Historically, many of these research communities have been largely unconnected. To bring them together to help forge a common understanding of principles, techniques, and application areas, we have held a number of cross-disciplinary conferences, and special issues of journals.

Conferences

In 2007 we inaugurated a new annual conference, the *International Conference on Computational Cultural Dynamics (ICCCD)*. After hosting this conference for three years (2007, 2008, and 2009) under the generous sponsorship of AFOSR, as well as under the banner of the Association for the Advancement of Artificial Intelligence (who published the proceedings), in 2010, we merged it with another conference that had been founded at about the same time by Huan Liu (of Arizona State University) and John Salerno (of AFRL): the *International Conference on Social Computing, Behavioral Modeling, and Prediction (SBP)*. The combined conference retains the acronym SBP, with “Behavioral” replaced by “Behavioral-Cultural”. The combined conference was first held was March 2011. The conference chairs were Dana Nau (one of the PIs on this MURI) and Sun-Ki Chai (a sociologist at the University of Hawaii), and

the program chairs were Huan Liu and John Salerno. Out of 85 submissions to the conference, 20 were accepted for oral presentation, and another 26 were accepted for poster presentation. This means that the acceptance rate for oral presentation was 24%, which means that SBP has become a very competitive, high-quality conference. The combined 2012 Conference chairs are Jay Yang and V.S. Subrahmanian, while the Program Chair is Nathan Bos.

In 2007, we also started the Scalable Uncertainty Management (SUM) conference which has been running annually since then. The conference focused on decision making under conditions of uncertainty. The Program Chairs of the first SUM conference were Henri Prade and V.S. Subrahmanian, with Didier Dubois as the General Chair.

Special Issues of Journals

IEEE Intelligent Systems. The March/April 2008 issue of *IEEE Intelligent Systems* was a special issue on Computational Cultural Dynamics, based on the best papers from the 2007 ICCCD conference. It was edited by Dana Nau (one of the PIs on this MURI) and Jonathan Wilkenfeld (Government and Politics Dept., University of Maryland).

ACM Transactions on Intelligent Systems (TIST). The October 2010 issue of *ACM TIST* was a special issue on AI in Social Computing and Cultural Modeling, co-edited by Dana Nau and Huan Liu. Another special of *ACM TIST* is forthcoming, based on the best papers from the 2011 SBP conference.

II. Scientific Achievements

We summarize the different scientific advances made during the CARA project below.

II.A CARA Cultural Reasoning Architecture

We proposed the CARA Cultural Reasoning Architecture and developed many of the components within it. Prior to CARA, virtually all research on group behavior collected data manually. For instance, the Minorities at Risk Organizational Behavior (MAROB) project collected data manually about over 100 terror groups worldwide, but because of the manual

collection (which required a small army of students to gather data), they were only able to collect data on a coarse-grained, annual basis. Moreover, data was always at least 5 years behind the current time.

The CARA architecture, in contrast, proposed a paradigm that:

- *Automatically collected data from open sources information;*
- *Developed algorithms to automatically analyze/mine the data, identifying conditions on “environmental” variables that were potentially predictive of various actions carried out by the groups being modeled;*
- *Develop methods to present the derived “stochastic opponent modeling agent (SOMA)” rules to analysts using both:*
 - *A social networking method called the SOMA Terror Organization Portal (STOP) and*
 - *A massive multi-player online game platform (not covered in this report as it was not funded by this project).*

Publications:

V.S. Subrahmanian, M. Albanese, V. Martinez, D. Nau, D. Reforgiato, G. Simari, A. Sliva and J. Wilkenfeld. CARA: A Cultural Adversarial Reasoning Architecture, *IEEE Intelligent Systems*, Vol. 22, Number 2, pages 12—16, March/April 2007.

V.S. Subrahmanian. Cultural Modeling in Real-time, *Science*, Vol. 317, Nr. 5844, pages 1509-1510, Sep. 14, 2007.

V.S. Subrahmanian and J. Dickerson. What can virtual worlds do for national security? *Science*. Vol. 326, pages 1201-1202, 27 Nov. 2009.

II.B The T-REX RDF Extraction Engine

As an important part of CARA, we developed T-REX (short for “The RDF EXtractor”) which can automatically scan news articles and extract RDF from it. RDF represents information in the form of “triples” (s, p, o) consisting of a subject s , a property p , and an object o . Being structured, RDF triples are highly amenable to computational analysis in comparison to document that are

unstructured. In addition, a set of RDF triples can be represented also as a graph, allowing sophisticated network analysis methods to be unleashed.

T-REX attempts to automatically extract RDF via a *training* phase, a *learning* phase, and an *operational* phase.

- In the *training* phase, T-REX is invoked with a training corpus of sentences reflecting the types of information we want to extract automatically. For instance, if we want to extract information about violent events, the training corpus consists of sentences about violent events. An annotation interface shows a human annotator a parse tree of a sentence and asks him to annotate nodes in the parse tree. These annotations allow the human annotator to specify the *(subject, property, object)* triples in the sentence and also allow the annotator to flag parts of speech, replacements, etc. that can be used.
- In the *learning phase*, T-REX takes the annotated corpus as input and automatically *learns extraction rules* from the annotated corpus. Intuitively, these learned extraction rules say things like “if the parse tree has a node XXX of type TTT and there is a child that belongs to the category of KILL verbs, then infer that node YYY is the subject, node ZZZ is the object, and the property is KILL.
- In the *operation phase*, T-REX applies the extracted rules learned in the previous phase to documents that it has never seen before.

The general principles and theory behind T-REX are very generic and can be used to extract a wide variety of information on many different domains. However, in the context of the CARA project, we applied T-REX to the automated extraction of violent events worldwide. T-REX’s Violent Event Extractor, when it was fully operational, could track over 100,000 articles a day from 93 countries around the world.

As part of the T-REX effort, we found that RDF was not adequate to express many aspects of real world events. For instance, RDF was not adequate to capture time (example time or duration of an event).

As a consequence, we proposed *Annotated RDF*, an extension of RDF consisting of quadruples (s, p, o, a) where (s, p, o) are as before and a is any member of an arbitrary, but fixed lattice \mathbf{L} . We

described the syntax of aRDF, the logical semantics of aRDF, and developed sophisticated algorithms to query aRDF data sets efficiently.

Publications:

O. Udrea, D. Reforgiato Recupero, and V.S. Subrahmanian. Annotated RDF, *ACM Transactions on Computational Logic*, Vol. 11, Nr. 2, pages 1-41, Jan. 2010, <http://doi.acm.org/10.1145/1656242.1656245>

M. Albanese and V.S. Subrahmanian. T-REX: A System for Automated Cultural Information Extraction, *Proc. 2007 International Conference on Computational Cultural Dynamics*, College Park, MD, pages 2-8, AAAI Press, Aug. 2007.

O. Udrea, D. Reforgiato and V.S. Subrahmanian. Annotated RDF, *Proc. 2006 European Semantic Web Conference*, Budva, Montenegro, Lecture Notes in Computer Science, Volume 4011, June 2006, Pages 487 - 501, DOI 10.1007/11762256_36, URL http://dx.doi.org/10.1007/11762256_36

[Octavian Udrea](#), V. S. Subrahmanian, [Zoran Majkic](#): Probabilistic RDF. Proc. IEEE [IRI 2006](#): 172-177.

II.C SOMA: Stochastic Opponent Modeling Agents

Stochastic Opponent Modeling Agents (SOMA for short) are sets of probabilistic rules that can be associated with any given terrorist group. Suppose we are given a group G and table T(G) containing information about the group in the following format:

- Rows of T(G) correspond to different periods of time;
- Columns of T(G) fall into two categories:
 - Columns reflecting *environmental attributes* describing the social, cultural, political, economic environment within which the group was operating and
 - Columns reflecting *action attributes* describing the intensity of different types of actions taken by the group.

Our SOMA framework takes a table of the type described above and *automatically* extracts SOMA rules from it. Soma-rules have the form

$$A(i):p \leftarrow B_1:p_1 \ \& \ \dots \ \& \ B_n:p_n$$

where A is an *action atom* (involving an action attribute) and the B_i's are an *environmental atoms* involving environmental attributes. Such rules can informally be read as: “If each of the

B_i 's is true about the group's environment in a given time period with certainty at least p_i , then the group will take action A at a given intensity i with certainty p ."

In our work, we:

- Formally defined the syntax of SOMA rules;
- Gave a formal possible-worlds based semantics for SOMA rules;
- Studied the computational complexity of SOMA rules;
- Showed that a set of SOMA rules is consistent iff a corresponding linear program can be solved;

In addition, we developed algorithms to solve the *Most Probable World* problem in which we try to identify the set of actions that the terrorist group will take, given a user-provided (or hypothetical) setting of the environmental variables. To find the most probable reaction to a given environmental setting, we set up a linear program which can be extremely large – in realistic applications, such a linear program can have anywhere between 2^{1000} to 2^{30000} variables, making it infeasible to even write it down, let alone solve it. Moreover, we need to solve an equally large set of such linear programs.

As this is clearly infeasible, we develop fast algorithms to solve such problems approximately by using a mix of four techniques:

- Randomization by selecting random sets of variables to consider;
- Tuned linear programs that allow the original linear program to be replaced by a much smaller and much simpler linear program;
- A sophisticated search technique that allows us to iteratively fine-tune the above linear program by solving it and tuning it again.
- Parallel algorithms to solve these linear programs very efficiently by generating multiple random samples and solving multiple linear programs in parallel.

Publications

G. Simari, V. Martinez, A. Sliva and V.S. Subrahmanian. Focused Most Probable World Computations in Probabilistic Logic Programs, accepted for publication in *Annals of Math and Artificial Intelligence*.

S. Khuller, V. Martinez, D. Nau, G. Simari, A. Sliva and V.S. Subrahmanian. Computing Most Probable Worlds of Action Probabilistic Logic Programs: Scalable Estimation for $10^{30,000}$ worlds, *Annals of Mathematics and Artificial Intelligence*, Vol. 51, No, 2-4, pages 295-331, Dec. 2007.

V.S. Subrahmanian. Cultural Modeling in Real-time, *Science*, Vol. 317, Nr. 5844, pages 1509-1510, Sep. 14, 2007.

G. I. Simari, V. Martinez, A. Sliva and V.S. Subrahmanian. Scaling Most Probable World Computations in Probabilistic Logic Programs, *Proc. 2008 Intl. Conf. on Scalable Uncertainty Management*, Oct. 2008, Naples, Italy. Springer Verlag Lectures Notes in Computer Science.

Aaron Mannes, Amy Sliva, V.S. Subrahmanian and Jonathan Wilkenfeld. Stochastic Opponent Modeling Agents: A Case Study with Hamas, *Proc. 2008 Intl. Conf. on Computational Cultural Dynamics*, pages 49-54, Sep. 2008, AAAI Press.

V. Martinez, G. Simari, A. Sliva and V.S. Subrahmanian. The SOMA Terror Organization Portal (STOP): Social Network and Analytic Tools for the Real-Time Analysis of Terror Groups, in *Proc. 2008 First Intl. Workshop on Social Computing, Behavioral Modeling and Prediction* (eds. H. Liu, J. Salerno and M. Young), pages 9-18, Phoenix, April 1-2, 2008, Springer Verlag.

A. Mannes, M. Michael, A. Pate, A. Sliva, V.S. Subrahmanian and J. Wilkenfeld. Stochastic Opponent Modelling Agents: A Case Study with Hezbollah, *Proc. 2008 First Intl. Workshop on Social Computing, Behavioral Modeling and Prediction*, (eds. H. Liu, J. Salerno, and M. Rogers), Springer Verlag, Phoenix, April 1-2, 2008.

S. Khuller, V. Martinez, D. Nau, A. Sliva, G.I. Simari and V.S. Subrahmanian. Finding Most Probable Worlds of Probabilistic Logic Programs, *Proc. 2007 International Conference on Scalable Uncertainty Management*, Lecture Notes in Computer Science Vol. 4772, pages 45-59, Springer Verlag.

A. Sliva, V. Martinez, G. Simari and V.S. Subrahmanian. SOMA Models of the Behaviors of Stakeholders in the Afghan Drug Economy: A Preliminary Report. *Proc. 2007 International Conference on Computational Cultural Dynamics*, College Park, MD, pages 78-86, . AAAI Press, Aug. 2007.

G. Simari, A. Sliva, D. Nau and V.S. Subrahmanian. A Stochastic Language for Modelling Opponent Agents. *Proc. Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 244-246, Hakodate, Japan, May 2006.

II.D Real-World Terror Group Analysis

SOMA models of terror groups have been used successfully in reasoning about several real-world terror groups: notably Hezbollah, Hamas and Lashkar-e-Taiba. In addition to terror groups, we also looked at SOMA models of stakeholders in the Afghan drug economy.

In a paper we published in April 2008, we derived SOMA-rules that predicted conditions under which Hezbollah does not take launch transnational attacks against civilians in Israel. In particular, these conditions stated that (in conjunction with certain other conditions), Hezbollah usually does not launch such transnational attacks during periods when there are ongoing elections in Lebanon. In November 2008, the Beirut Daily Star published an article reviewing our paper with skeptical comments from both the reporter involved, a Lebanese political science professor, and a spokesman for Hezbollah. Our prediction showed that it was very likely that Hezbollah would not carry out transnational attacks on Israel in the first half of 2009 when these conditions were prevailing. Our prediction was proven correct. Thus, even though Hezbollah knew about our prediction, they still did what we predicted they would do because this is what made the most sense under the circumstances.

Publications:

A. Sliva, V. Martinez, G. Simari and V.S. Subrahmanian. SOMA Models of the Behaviors of Stakeholders in the Afghan Drug Economy: A Preliminary Report. *Proc. 2007 International*

Conference on Computational Cultural Dynamics, College Park, MD, pages 78-86, . AAAI Press, Aug. 2007.

Aaron Mannes, Amy Sliva, V.S. Subrahmanian and Jonathan Wilkenfeld. Stochastic Opponent Modeling Agents: A Case Study with Hamas, *Proc. 2008 Intl. Conf. on Computational Cultural Dynamics*, pages 49-54, Sep. 2008, AAAI Press.

A. Mannes, M. Michael, A. Pate, A. Sliva, V.S. Subrahmanian and J. Wilkenfeld. Stochastic Opponent Modelling Agents: A Case Study with Hezbollah, *Proc. 2008 First Intl. Workshop on Social Computing, Behavioral Modeling and Prediction*, (eds. H. Liu, J. Salerno, and M. Rogers), Springer Verlag, Phoenix, April 1-2, 2008.

A. Mannes and V.S. Subrahmanian. Calculated terror, *Foreign Policy magazine (online edition)*, Dec. 15, 2009 edition.

II.E Systems Social Science

This work by our subcontractor, the University of Pennsylvania, focuses on the following important question:

How can we best represent the numerous social science theories and models in the literature within a computational environment so that users, analysts, and others can evaluate the utility of these methods for this mission?

Systems social science is the attempt to synthesize esoteric (social science theories) with exoteric (SME mental models and stakeholder perspectives) knowledge in seeking to model social dilemmas and potential resolutions. The investigator has pursued systems social science for over a decade and has accumulated a model base of hundreds of social science theories ranging across psychology, anthropology, sociology, political science, and economics. These are organized into a library of models covering agent cognition (perception, physiology/stress, affect/values/personality, and decision style/coping mechanisms) and social worlds (relations/mobilization, group membership/culture, political institutions, economic sectors). The result is one of the few meso-scale socio-cognitive model libraries currently available, and thus an excellent model base to investigate for performance, adaptability, and transparency.

This model base has an excellent pedigree and performance record. The criterion for selection of models is comparable to the scale used to judge models in other sciences, and an attempt has been made to select best-of-breed theories for inclusion, though the library architecture has been established solely for the purposes of discarding current models and adding new models as the theory advances and better practices are discovered (or based on user preference). Further, the

entire collection has been subjected to what is perhaps the largest correspondence test in the HSCB modeling field: 320 specific forecasts were attempted in a 2008 challenge grant. In particular, five metrics (rebellion, insurgency, domestic violence, political crisis, and repression) were forecast in each of four Asian countries for each quarter of 2004-2007. These 5x4x16 quarterly model forecasts were compared to ground truth assembled by a third party and were found to have 80% accuracy on average.

In game theory, complexity sciences, and social agents, one often designs according to the KISS principle (Keep It Simple, Stupid). By contrast, in the white box, bottom up esoteric and exoteric synthesis approach taken here the KIDS principle (Keep It Descriptive, Stupid) applies, and agents and organizations are profiled in great detail – individualized moralistic payoff functions, diverse agendas and action portfolios, etc. Thus, in addition to the quarterly metrics, a great many more things were forecast such as the reasoning of all the agents in all these countries and the behaviors of the diverse political groups, institutions, and factions. These detailed outcomes are available for inspection in an effort to try and isolate causes of rebellions, insurgencies, and so on. Yet such inspections are daunting – not just because of the scale of the model library, but also because of inherent auto-correlation of starting conditions and potential hyper-confluence of the model parameter outputs both within and between the teleologic agents and organizations in the countries being simulated. And while the library of models has passed significant tests of performance and adaptability, it risks being seen as a grey box, at best, if the output data is impenetrable and causal connections indiscernible to users. As a result, we are launching new efforts to invest in and investigate the value of an array of methods to improve the model library in terms of input-to-output transparency, model replaceability and maintainability, traceability of outcomes back to likely causes including agent self-explainability, and progressive design inquiry and computational experimentation.

Technologically, the main product of our work has been the development of FactionSim on top of PMFserv, introducing a social layer to the existing cognitive framework. This adds the ability to model such things as groups, economies, and political institutions, which has led to four new applications being started:

- NonKin Village (a village-level training program for USMC/ONR)
- Human Terrain Tool (a village-level simulator for USMC/ONR)
- CountrySim (a country-level ethno-political conflict simulator for DARPA) and
- ICS (an international crisis simulator for DARPA).

Additionally, the software is being used at ARL (Aberdeen Proving Grounds) and JHU APL.

This work has led to a large number of publications, including two manuals, two public tutorials per year (on average), three journal articles, three (pending) book chapters, two public tutorials per year (on average), 18 conference or workshop papers, and 22 invited or keynote talks (on four continents). Some of these citations are included below. Finally, our work has resulted in two awards: the International Dissertation of the Year, INCOSE'06 and the Wharton Dissertation Award (Decision & Risk Center).

II.F Evolution of State-Dependent Risk Preferences

There is much empirical evidence that human decision-making under risk does not coincide with expected value maximization, and much effort has been invested into the development of descriptive theories of human decision-making involving risk (e.g. Prospect Theory). An open question is how behavior corresponding to these descriptive models could have been learned or arisen evolutionarily, as the described behavior differs from expected value maximization. We believe that one part of the answer is the interplay between risk-taking and sequentiality of choices in populations subject to evolutionary population dynamics.

To test this hypothesis, we have performed analyses and simulations on several evolutionary game simulations designed to study the risk behavior of agents in evolutionary environments. These include several evolutionary lottery games where sequential decisions are made between risky and safe choices. For the evolutionary population dynamic, we have considered all *imitation dynamics*—a class of population dynamics that includes the replicator dynamic (which has been widely used to model the biological evolution of animal populations), the imitate-the-better dynamic (which is thought to be a good model of how humans learn by observing the behavior of others), and a range of population dynamics intermediate between them.

Our results demonstrate that due to the interplay between risk-taking and sequentiality of choices, there is a state-dependent strategy—namely, to *take risks when one is behind, and play it safe when one is ahead*—that has an evolutionary advantage over most other strategies, including the well-known strategy of expected-value maximization; and the evolutionary advantage occurs with *every* imitation dynamic except for the replicator dynamic.

The above strategy embodies a risk preference similar to several prominent models of human decision making, as well as some well-known practical examples of human decision making—for example, the “Hail Mary play” in American football, and analogous behavior various other situations. This suggests that population dynamics other than the replicator dynamic may model an important mechanism for the emergence of those risk preferences.

We have further shown, in an evolutionary version of the well-known stag hunt game, how the state-dependent strategy described above can facilitate the evolution of cooperation in situations where cooperation entails risk.

Publications:

- P. Roos and D. S. Nau. Conditionally risky behavior vs expected value maximization in evolutionary games. In Sixth Conference of the European Social Simulation Association (ESSA 2009), Sept. 2009.
- P. Roos and D. S. Nau. State-dependent risk preferences in evolutionary games. In S.-K. Chai, J. J. Salerno, and P. L. Mabry, editors, *Advances in Social Computing: Third International Conference on Social Computing, Behavioral Modeling, and Prediction, SBP 2010*, volume LNCS 6007, pp. 23–31. Springer, Mar. 2010.
- P. Roos and D. S. Nau. Risk preference and sequential choice in evolutionary games. *Advances in Complex Systems* 13(4):559–578, Aug. 2010.
- P. Roos, J. R. Carr, and D. S. Nau. Evolution of state-dependent risk preferences. *ACM Transactions on Intelligent Systems and Technology (TIST)* 1(1):6:1–6:21, Oct. 2010.

II.G Evolution of Cooperative Societies

We have used evolutionary game theory to study the evolution of cooperative societies and the behaviors of individual agents (i.e., players) in such societies. We have developed a novel player model based upon empirical evidence from the social and behavioral sciences stating that: (1) an individual’s behavior may often be motivated not only by self-interest but also by the consequences for others, and (2) individuals vary in their interpersonal social tendencies, which reflect stable personal orientations that influence their choices.

Alongside the formal player model we have provided an analysis that considers possible interactions between different types of individuals and identifies five general steady-state behavioral patterns. We have developed evolutionary simulations that ratify previous findings on the evolution of cooperation, and provide new insights on the evolutionary process of cooperative behavior in a society as well as on the emergence of cooperative societies. Our main experimental result demonstrates that in contrast to previous common knowledge, increasing mutual reward or mutual punishment in the Prisoner’s Dilemma game does not result in the same type of cooperative society: while increasing reward does increase the society’s cooperativeness level, increasing mutual punishment does not.

Publications

- K.-L. Cheng, I. Zuckerman, U. Kuter, and D. S. Nau. Emergence of cooperative societies in evolutionary games. In Fourth Evolutionary Computation and Multi-Agent Systems and Simulation Workshop (ECoMASS'10), pp. 1794–1800, July 2010.
- K.-L. Cheng, I. Zuckerman, U. Kuter, and D. S. Nau. Evolving cooperative societies. In Late Breaking abstracts workshop, Genetic and Evolutionary Computation Conference (GECCO'10), pp. 2067–2068, July 2010.
- K.-L. Cheng, I. Zuckerman, U. Kuter, and D. Nau. Using a social orientation model for the evolution of cooperative societies. In 2010 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT-10), Sept. 2010.
- B. Wilson, I. Zuckerman, and D. S. Nau. Modeling social preferences in multi-player games. In International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2011.

II.H Distinguishing Intentional Actions from Noise

The Iterated Prisoner's Dilemma (IPD) has become well known as an abstract model of a class of multi-agent environments in which agents accumulate payoffs that depend on how successful they are in their repeated interactions with other agents. An important variant of the IPD is the Noisy IPD, in which there is a small probability, called the *noise level*, that accidents will occur. In other words, the noise level is the probability of executing “cooperate” when “defect” was the intended move, or vice versa.

Accidents can cause difficulty in maintaining cooperations with others in real-life situations, and the same is true in the Noisy IPD. Strategies that do quite well in the ordinary (non-noisy) IPD may do quite badly in the Noisy IPD. For example, if two players both use the well-known Tit-For-Tat (TFT) strategy, then an accidental defection may cause a long series of defections by both players as each of them punishes the other for defecting.

We have developed a technique called *symbolic noise detection*, for detecting whether anomalies in player's behavior are deliberate or accidental. The key idea is to construct a model of the other agent's behavior, and watch for any deviation from this model. If the other agent's next action is inconsistent with this model, the inconsistency can be due either to noise or to a genuine change

in their behavior; and we can often distinguish between two cases by waiting to see whether this inconsistency persists in next few moves.

We entered several different versions of our strategy in the 20th *Anniversary Iterated Prisoner's Dilemma Competition*, in Category 2 (noisy environments). Out of the 165 contestants in this category, our programs consistently ranked among top ten. The best of our programs ranked third, and it was beaten only by two “master-slave strategy” programs that each had a large number of “slave” programs feeding points to them.

Publications:

- T.-C. Au, U. Kuter, and D. S. Nau. Planning for interactions among autonomous agents. In K. Hindriks, A. Pokahr, and S. Sardina, editors, *ProMAS'08 Post-Proceedings*. Springer, 2009. Invited paper.
- T.-C. Au, D. S. Nau, and S. Kraus. Synthesis of strategies from interaction traces. In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 855–862, May 2008.
- T.-C. Au, S. Kraus, and D. S. Nau. Symbolic noise detection in the noisy iterated chicken game and the noisy iterated battle of the sexes. In *International Conference on Computational Cultural Dynamics (ICCCD)*, Aug. 2007.
- T.-C. Au and D. S. Nau. Reactive query policies: A formalism for planning with volatile external information. In *IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, pp. 243–250, 2007.
- T.-C. Au and D. S. Nau. Is it accidental or intentional? a symbolic approach to the noisy iterated prisoner's dilemma. In G. Kendall, X. Yao, and S. Y. Chong, editors, *The Iterated Prisoners Dilemma: 20 Years On*, pp. 231–262. World Scientific, 2007.
- T.-C. Au and D. S. Nau. The incompleteness of planning with volatile external information. In *European Conference on Artificial Intelligence (ECAI)*, pp. 839–840, Aug. 2006.
- T.-C. Au and D. S. Nau. Maintaining cooperation in noisy environments. In *National Conference on Artificial Intelligence (AAAI)*, pp. 1561–1564, July 2006.

- T.-C. Au and D. S. Nau. Accident or intention: That is the question (in the iterated prisoner's dilemma). In International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 561–568, 2006.

II.I Social Learning

Social learning, i.e., learning from observations of the behaviors of others and the outcomes of those behaviors, is an important capability in many animal species. It is particularly important in humans, and some social learning theorists believe it is how most human behavior is learned. There is evidence that animals and humans do not indiscriminately copy the behaviors of others, but instead use strategies that involve evaluating the payoffs that others receive for their behaviors. However, much is unknown about what strategies work best and how they might have evolved. For example, it seems natural to assume that communication has evolved due to the inherent superiority of copying others' success rather than learning on one's own via trial-and-error innovation. However, there has also been substantial work questioning this intuition. Several evolutionary games have been developed to investigate social learning. One of the best-known is Cultaptation, a multi-agent social-learning game developed by a consortium of European scientists. The same consortium sponsored an international tournament that carried a €10,000 prize.

We have produced an extensive set of mathematical and simulation results for Cultaptation. Our work has had two main objectives: (1) to study the nature of Cultaptation to see what types of strategies are effective; and (2) more generally, to develop ways of analyzing evolutionary environments with social learning. Our results provide strong evidence for the following hypotheses:

1. The best strategies for Cultaptation and similar games are likely to be conditional ones in which the choice of action at each round is conditioned on the agent's accumulated experience.
2. Such strategies (or close approximations of them) can be computed by doing a lookahead search that predicts how each possible choice of action at the current round is likely to

affect future performance.

3. Such strategies are likely to exploit most of the time, but will have ways of quickly detecting structural shocks, so that they can switch quickly to innovation in order to learn how to respond to such shocks. This conflicts with the conventional wisdom that successful social-learning strategies are characterized by a high frequency of innovation; and agrees with recent experiments by others on human subjects that also challenge the conventional wisdom.

Publications

- R. Carr, E. Raboin, A. Parker, and D. S. Nau. When innovation matters: An analysis of innovation in a social learning game. In International Conference on Computational Cultural Dynamics (ICCCD), Sept. 2008.
- E. Raboin, R. Carr, A. Parker, and D. S. Nau. Balancing innovation and exploitation in a social learning game. In AAAI Fall Symposium on Adaptive Agents in Cultural Contexts, Nov. 2008.
- R. Carr, E. Raboin, A. Parker, and D. Nau. Within epsilon of optimal play in the cultaptation social learning game. In International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2009. Extended abstract.
- R. Carr, E. Raboin, and D. Nau. Generation and Analysis of Strategies in an Evolutionary Social Learning Game. In preparation for journal submission, 2012.

Technical Transitions

An important part of our work is to transition our research to both the DoD and to the commercial marketplace. We were fortunate to have several important transitions:

- The SOMA Terror Organization Portal, developed under this effort, had registered users from over 15 national security related agencies. Registered users included individuals from NSA, CIA, DIA, Lackland AFB, Rome Labs, Dept. of State, and JWAC, among others.
- In a related effort, the SOMA paradigm was used to develop certain very specific models of three tribes for the US Army TRADOC Intelligence Support Activity (TRISA). These models were sent to TRISA and to the US Army AMSAA agencies.

- The T-REX violence watch program developed in part under this contract has had registered users from 4 national security related agencies.

IV. Publications

Archival publications published during the reporting period:

1. V.S. Subrahmanian and A. Kruglanski , Editors, Proceedings of the 2nd International Conference on Computational Cultural Dynamics, AAAI Press, Sep. 2008. ISBN 978-1-57735-389-8 (book), 978-1-57735-390-4 (CD).
2. V.S. Subrahmanian, Editor, Special Issue of Annals of Mathematics and Artificial Intelligence Journal celebrating Jack Minker's 80th birthday, Vol. 51, Nos. 2—4, Dec. 2007.
3. V.S. Subrahmanian, Editor (with H. Prade). Proceedings – 2007 International Conference on Scalable Uncertainty Management, Springer-Verlag Lecture Notes in Computer Science, Volume 4772, October 2007, Springer-Verlag, Heidelberg.
4. A. Sliva, V.S. Subrahmanian, V. Martinez and G. Simari. CAPE: Automatically Predicting Changes in Terror Group Behavior, accepted for publication in *Mathematical Methods in Counterterrorism* (eds. N. Memon), Springer Verlag, to appear.
5. V.S. Subrahmanian. Probabilistic Temporal Databases, in *Encyclopedia of Temporal Databases* (eds. T. Ozsu and L. Liu), Springer Verlag, accepted, to appear, 2008.
6. V.S. Subrahmanian, V. Martinez and D. Reforgiato. Multimedia Presentation Databases, in *Encyclopedia of Temporal Databases* (eds. T. Ozsu and L. Liu), Springer Verlag, under review, 2008.
7. R. Chellappa, N.P. Cuntoor, S.W. Joo, V.S. Subrahmanian and P. Turaga. Computational Vision Approaches for Event Modeling. In "Understanding Events: From Perception to Action", T. F. Shipley and J. Zacks (eds.), pages 473-521, Oxford University Press, 2008.
8. A. Parker, D. Nau and V.S. Subrahmanian. The Role of Imperfect Information, in "Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind" (eds. A. Kott and W.M. McEneaney), pages 209—229, Chapman and Hall, 2006.
9. M. Albanese, R. Chellappa, V. Moscato, A. Picariello, V.S. Subrahmanian P. Turaga and O. Udrea. A Constrained Probabilistic Petri Net Model for Human Activity Detection in Video, *IEEE Transactions on Multimedia*, Vol. 10, Nr. 6, pages 982-996, Oct. 2008.
http://ieeexplore.ieee.org/xpls/abs_all.jsp?isnumber=4657444&arnumber=4657459&count=29&index=5

10. O. Udrea, D. Reforgiato Recupero, and V.S. Subrahmanian. Annotated RDF, *ACM Transactions on Computational Logic*, Vol. 11, Nr. 2, pages 1-41, Jan. 2010, <http://doi.acm.org/10.1145/1656242.1656245>
11. R. Chellappa, P. Turaga, V.S. Subrahmanian, and O. Udrea. Machine Recognition of Human Activities: A Survey, *IEEE Transactions on Circuit for Systems and Video Technology*, Vol. 18, Nr. 11, pages 1473-1488, Nov. 2008.
12. J. Grant, G. Infantes, A. Parker and V.S. Subrahmanian. SPOT Databases: Efficient Consistency Checking and Optimistic Selection in Probabilistic Spatial Databases, accepted for publication in *IEEE Transactions on Knowledge and Data Engineering*, April 2008.
13. V.Martinez, G. Simari, A. Sliva and V.S. Subrahmanian. CONVEX: Similarity-based Algorithms for Forecasting Group Behavior, *IEEE Intelligent Systems*, Vol. 23, 4, pages 51-57, July/Aug 2008.
14. S. Khuller, V. Martinez, D. Nau, G. Simari, A. Sliva and V.S. Subrahmanian. Computing Most Probable Worlds of Action Probabilistic Logic Programs: Scalable Estimation for $10^{30,000}$ worlds, *Annals of Mathematics and Artificial Intelligence*, Vol. 51, No. 2-4, pages 295-331, Dec. 2007.
15. Y. Zhang, E. Manister, S. Kraus, V.S. Subrahmanian, D. Peleg. Computing the survivability of a multi-agent deployment, *Artificial Intelligence journal*.
16. V.S. Subrahmanian. Cultural Modeling in Real-time, *Science*, Vol. 317, Nr. 5844, pages 1509-1510, Sep. 14, 2007.
17. A. Parker, J. Grant, and V.S. Subrahmanian. A Logical Formulation of Probabilistic Spatial Databases, *IEEE Transactions on Knowledge and Data Engineering*, Vol. 19, Nr. 11, pages 1541-1556.
18. M. Fayzullin, V.S. Subrahmanian, M. Albanese, C. Cesarano, A. Picariello. Story Creation from Heterogeneous Data Sources, *Multimedia Tools and Applications*, Vol. 33, Nr. 3, pages 351-377, June 2007.
19. V.S. Subrahmanian, M. Albanese, V. Martinez, D. Nau, D. Reforgiato, G. Simari, A. Sliva and J. Wilkenfeld. CARA: A Cultural Adversarial Reasoning Architecture, *IEEE Intelligent Systems*, Vol. 22, Number 2, pages 12—16, March/April 2007. Available at http://www.computer.org/portal/cms_docs_intelligent/intelligent/content/Promo/x2012_07.pdf
20. M. Fayzullin, V.S. Subrahmanian, M. Albanese, C. Cesarano and A. Picariello. The Priority Curve Algorithm for Video Summarization, *Information Science*, Vol. 31, pages 679-695, 2006.
21. E. Hung, L. Getoor and V.S. Subrahmanian. Probabilistic Interval XML, *ACM Transactions on Computational Logic*, Vol. 8, Nr. 4, August 2007.

22. G. I. Simari, V. Martinez, A. Sliva and V.S. Subrahmanian. Scaling Most Probable World Computations in Probabilistic Logic Programs, *Proc. 2008 Intl. Conf. on Scalable Uncertainty Management*, Oct. 2008, Naples, Italy. Springer Verlag Lectures Notes in Computer Science.
23. Aaron Mannes, Amy Sliva, V.S. Subrahmanian and Jonathan Wilkenfeld. Stochastic Opponent Modeling Agents: A Case Study with Hamas, *Proc. 2008 Intl. Conf. on Computational Cultural Dynamics*, pages 49-54, Sep. 2008, AAAI Press.
24. John Dickerson, Maria Vanina Martinez, Diego Reforgiato and V.S. Subrahmanian, University of Maryland. CIG: Cultural Island and Games, *Proc. 2008 Intl. Conf. on Computational Cultural Dynamics*, pages 26-31, Sep. 2008, AAAI Press.
25. V. Martinez, F. Parisi, A. Pugliese, G. Simari and V.S. Subrahmanian. Inconsistency Management Policies, *Proc. KR 2008*, Sydney, Australia, Sep. 2008
26. G. Simari, M. Broecheler, V.S. Subrahmanian and Sarit Kraus. Promises Kept, Promises Broken: An Axiomatic and Quantitative Treatment of Fulfilment, *Proc. KR 2008*, Sydney, Australia, Sep. 2008.
27. A. Parker, G. Infantes, J. Grant and V.S. Subrahmanian. An AGM-based Belief Revision Mechanism for Probabilistic Spatio-Temporal Logics, *Proc. AAAI 2008*, pages 606-614, Boston, MA. ***One of a relatively small number of papers at AAAI 2008 accepted for both oral and poster presentation.***
28. A. Pugliese, O. Udrea and V.S. Subrahmanian. Scalable Indexing for Temporal RDF, *Proc. 2008 International World Wide Web Conference (WWW '08)*, Beijing, China, pages 605-614, April 2008. <http://doi.acm.org/10.1145/1367497.1367579>
29. V. Martinez, G. Simari, A. Sliva and V.S. Subrahmanian. The SOMA Terror Organization Portal (STOP): Social Network and Analytic Tools for the Real-Time Analysis of Terror Groups, in *Proc. 2008 First Intl. Workshop on Social Computing, Behavioral Modeling and Prediction* (eds. H. Liu, J. Salerno and M. Young), pages 9-18, Phoenix, April 1-2, 2008, Springer Verlag.
30. A. Mannes, M. Michael, A. Pate, A. Sliva, V.S. Subrahmanian and J. Wilkenfeld. Stochastic Opponent Modelling Agents: A Case Study with Hezbollah, *Proc. 2008 First Intl. Workshop on Social Computing, Behavioral Modeling and Prediction*, (eds. H. Liu, J. Salerno, and M. Rogers), Springer Verlag, Phoenix, April 1-2, 2008.
31. V. Martinez, A. Pugliese, G. I. Simari, V.S. Subrahmanian and H. Prade. How Dirty is your Relational Database?, *Proc. ECSQARU 2007*, pages 103-114.

32. M. Albanese, V. Moscato, A. Picariello, V.S. Subrahmanian, O. Udrea. Detecting Stochastically Defined Activities in Video, Proc. IJCAI 2007, pages 1802—1807, Jan. 2007, Hyderabad, India.
33. T.-C. Au, U. Kuter, and D. S. Nau. Planning for interactions among autonomous agents. In International Workshop on Programming Multi-Agent Systems (ProMAS), 2008, Invited paper, to appear.
34. T.-C. Au, U. Kuter, and D. Nau. Web service composition with volatile information. In Proceedings of the International Semantic Web Conference (ISWC), 2005.
35. T.-C. Au and D. Nau. Accident or intention: That is the question (in the iterated prisoner's dilemma). In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2006.
36. T.-C. Au and D. Nau. Is it accidental or intentional? a symbolic approach to the noisy iterated prisoner's dilemma. In G. Kendall, X. Yao, and S. Y. Chong, editors, *The Iterated Prisoners Dilemma: 20 Years On*, pp. 231–262. World Scientific, 2007.
37. T.-C. Au and D. Nau. Reactive query policies: A formalism for planning with volatile external information. In *IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, pp. 243–250, 2007.
38. T.-C. Au, D. Nau, and S. Kraus. Synthesis of strategies from interaction traces. In International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2008.
39. F. Benamara, C. Cesarano, A. Picariello, D. Reforgiato, V.S. Subrahmanian. Sentiment Analysis: Adverbs and Adjectives are Better than Adjectives Alone, Proc. 2007 Intl. Conf. on the Web and Social Media, pages 203—206, Boulder, CO, March 2007.
40. R. Carr, E. Raboin, A. Parker, and D. Nau. Balancing innovation and exploitation in a social learning game. In AAI Fall Symposium on Adaptive Agents in Cultural Contexts, Nov. 2008, To appear.
41. R. Carr, E. Raboin, A. Parker, and D. Nau. When innovation matters: An analysis of innovation in a social learning game. In International Conference on Computational Cultural Dynamics (ICCCD), Sept. 2008.
42. C. Cesarano, A. Picariello, D. Reforgiato, V.S. Subrahmanian. OASYS 2: An Opinion Analysis System, Proc. 2007 Intl. Conf. on the Web and Social Media, pages 313-314, Boulder, CO, March 2007.
43. S. Khuller, V. Martinez, D. Nau, G. Simari, A. Sliva, and V. Subrahmanian. Finding most probable worlds of probabilistic logic programs. In International Conference on Scalable Uncertainty Management (SUM 2007), Oct. 2007.

44. U. Kuter, D. Nau, M. Pistore, and P. Traverso. Task decomposition on abstract states, for planning under nondeterminism. *Artificial Intelligence*, 2008, To appear.
45. U. Kuter, D. Nau, E. Reisner, and R. Goldman. Conditionalization: Adapting forward-chaining planners to partially observable environments. In *ICAPS 07 Workshop on Planning and Execution for Real-World Systems*, Sept. 2007.
46. U. Kuter, D. Nau, E. Reisner, and R. Goldman. Using classical planners to solve nondeterministic planning problems. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, Sept. 2008.
47. D. Nau. Current trends in automated planning. *AI Magazine* 28(4):43–58, 2007.
48. A. Parker, F. Yaman, D. S. Nau, and V. S. Subrahmanian. Probabilistic go theories. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 501–506, 2007.
49. V.S. Subrahmanian and L. Amgoud. A General Framework for Reasoning about Inconsistency, *Proc. IJCAI 2007*, pages 599-604, Jan. 2007, Hyderabad, India.
50. Silverman, BG, Bharathy, GK, Eidelson, RJ, Smith, TE, Nye, B. Sociocultural Games for Training and Analysis. *IEEE Transactions on Systems, Man and Cybernetics, Part A*, 37(6):1113-1130, 2007.
51. Silverman, BG, Bharathy, GK, Nye, B, Eidelson, RJ. Modeling Factions for “Effects Based Operations”: Part I – Leaders and Followers. *Computational & Mathematical Organization Theory*, 13(4):379-406, 2007.
52. Silverman, BG, Bharathy, GK, Nye, B, Smith, TE. Modeling Factions for “Effects Based Operations”: Part II – Behavioral Game Theory. *Computational & Mathematical Organization Theory*, 14(2):120-155, 2008.
53. Silverman, BG, Bharathy, GK, Kim, GJ. Challenges of Country Modeling with Databases, Newsfeeds, and Expert Surveys. In A Uhrmacher and D Weyns, editors, *Agents, Simulation and Applications*. Taylor and Francis, to appear.
54. Silverman, BG, Bharathy, GK, Nye, B. Gaming and Simulating Ethno-political Conflicts. In *Proceedings of the Descartes Conference on Mathematical Modeling of Counter-Terrorism*. Springer, to appear.
55. Silverman, BG, Bharathy, GK, Nye, B, Kim, GJ, Roddy, M, Poe, M. State and Sub-State Actors: Synthesizing Theories Across the Social Sciences. In A Kott, editor, *Estimating Impact: A Handbook of Computational Methods and Models for Anticipating Economic, Social, Political and Security Effects in International Interventions*. Springer, to appear.

56. V.S. Subrahmanian and J. Dickerson. What can virtual worlds do for national security?, *Science*. Vol. 326, pages 1201-1202, 27 Nov. 2009.
57. M. Albanese and V.S. Subrahmanian. T-REX: A System for Automated Cultural Information Extraction, *Proc. 2007 International Conference on Computational Cultural Dynamics*, College Park, MD, pages 2-8, AAAI Press, Aug. 2007.
58. O. Udrea, D. Reforgiato and V.S. Subrahmanian. Annotated RDF, *Proc. 2006 European Semantic Web Conference*, Budva, Montenegro, Lecture Notes in Computer Science, Vol. 4011, June 2006, Pages 487-501, http://dx.doi.org/10.1007/11762256_36
59. O. Udrea, V.S. Subrahmanian, Zoran Majkic. Probabilistic RDF. *Proc. IEEE IRI 2006*: 172-177.
60. G. Simari, V. Martinez, A. Sliva and V.S. Subrahmanian. Focused Most Probable World Computations in Probabilistic Logic Programs, accepted for publication in *Annals of Math and Artificial Intelligence*.
61. A. Sliva, V. Martinez, G. Simari and V.S. Subrahmanian. SOMA Models of the Behaviors of Stakeholders in the Afghan Drug Economy: A Preliminary Report. *Proc. 2007 International Conference on Computational Cultural Dynamics*, College Park, MD, pages 78-86,. AAAI Press, Aug. 2007.
62. G. Simari, A. Sliva, D. Nau and V.S. Subrahmanian. A Stochastic Language for Modelling Opponent Agents. *Proc. Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 244-246, Hakodate, Japan, May 2006.
63. A. Mannes and V.S. Subrahmanian. Calculated terror, *Foreign Policy magazine (online edition)*, Dec. 15, 2009 edition.
64. P. Roos and D.S. Nau. Conditionally risky behavior vs expected value maximization in evolutionary games. In Sixth Conference of the European Social Simulation Association (ESSA 2009), Sept. 2009.
65. P. Roos and D.S. Nau. State-dependent risk preferences in evolutionary games. In S.-K. Chai, J.J. Salerno, and P.L. Mabry, editors, *Advances in Social Computing: Third International Conference on Social Computing, Behavioral Modeling, and Prediction*, SBP 2010, volume LNCS 6007, pp. 23-31. Spring, Mar. 2010.
66. P. Roos and D.S. Nau. Risk preference and sequential choice in evolutionary games. *Advances in Complex Systems* 13(4):559-578, Aug. 2010.
67. P. Roos, J.R. Carr and D.S. Nau. Evolution of state-dependent risk preferences. *ACM Transactions on Intelligent Systems and Technology (TIST)* 1(1):6:1-6:21, Oct. 2010.

68. K.-L. Cheng, I. Zuckerman, U. Kuter and D.S. Nau. Emergence of cooperative societies in evolutionary games. In Fourth Evolutionary Computation and Multi-Agent Systems and Simulation Workshop (ECoMASS'10), pp. 1794-1800, July 2010.
69. K.-L. Cheng, I. Zuckerman, U. Kuter and D.S. Nau. Evolving cooperative societies. In Late Breaking abstracts workshop, Genetic and Evolutionary Computation Conference (GECCO'10), pp. 2067-2068, July 2010.
70. K.-L. Cheng, I. Zuckerman, U. Kuter and D.S. Nau. Using a social orientation model for the evolution of cooperative societies. In 2010 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT-10), Sept. 2010.
71. B. Wilson, I. Zuckerman and D.S. Nau. Modeling social preferences in multi-player games. In International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2011.
72. T.-C. Au, S. Kraus and D.S. Nau. Symbolic noise detection in the noisy iterated chicken game and the noisy iterated battle of the sexes. In International Conference on Computational Cultural Dynamics (ICCCD), Aug. 2007.
73. T.-C. Au and D.S. Nau. The incompleteness of planning with volatile external information. In European Conference on Artificial Intelligence (ECAI), pp. 839-840, Aug. 2006.
74. T.-C. Au and D.S. Nau. Maintaining cooperation in noisy environments. In National Conference on Artificial Intelligence (AAAI), pp. 1561-1564, July 2006.
75. R. Carr, E. Raboin, A. Parker and D.S. Nau. Within epsilon of optimal play in the cultaptation social learning game. In International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2009. Extended abstract.
76. R. Carr, E. Raboin and D.S. Nau. Generation and Analysis of Strategies in an Evolutionary Social Learning Game. In preparation for journal submission, 2012.